

ORIGINAL

DOCKET FILE COPY ORIGINAL

EX PARTE OR LATE FILED



Richard N. Clarke

Room 5462C2
295 North Maple Avenue
Basking Ridge, NJ 07920
908-221-8685

December 24, 1997

RECEIVED

DEC 24 1997

FEDERAL COMMUNICATIONS COMMISSION
OFFICE OF THE SECRETARY

Ms. Magalie Roman Salas
Secretary
Federal Communications Commission
1919 M. St., NW, Room 222
Washington, D.C. 20554

RE: Ex Parte Presentation - Proxy Cost Models
CC Docket No. 96-45

Dear Ms. Salas:

On December 23, 1997, AT&T and MCI (the Hatfield Model Sponsors or "HMS") met with Brian Clopton, Chuck Keller, Bob Loube, Richard Smith and Natalie Wales of the Universal Service Branch of the Common Carrier Bureau in regards to the staff's examination of cost models for universal service in CC Docket Nos. 96-45 and 97-160. The HMS were represented by Richard Clarke and Mike Lieberman of AT&T, Chris Frentrup of MCI, Chris Antis of PNR, and Brian Pitkin of Klick, Kent and Allen.

The attached document detailing how the Hatfield Model, v5.0 compares with the BCPM3 on the issues of customer counting, location and clustering was provided to the Commission staff at this meeting. Unfortunately, only the figures from this document were included in the ex parte filing that was made by the HMS on December 23, 1997. Please add the text of this document to the record, and associate it with the figures that were filed yesterday.

Two copies of this Notice are being submitted to the Secretary of the FCC in accordance with Section 1.1206(a)(2) of the Commission's rules.

Sincerely,

Richard N. Clarke /hc

Richard N. Clarke

Attachment

cc: Sheryl Todd	Natalie Wales
Brian Clopton	Brad Wimmer
Chuck Keller	Mark Kennet
Bob Loube	William Sharkey
Richard Smith	Anthony Bush

12/24/97 045
12/24/97

RECEIVED

DEC 24 1997

FEDERAL COMMUNICATIONS COMMISSION
OFFICE OF THE SECRETARY

Scorecard

Modeling Customer Location:

Hatfield 5.0 vs. BCPM3

AT&T and MCI

December 23, 1997

1. Unit of Measure

Issue: What unit of measure should be used to count customer locations for universal service purposes. The major alternatives are (in order of count size): total housing units (inhabited plus uninhabited), total households (equivalent to housing units that are inhabited), or households with telephones. Because uninhabited housing units, or households without telephones tend to be located in more remote areas than households with telephones, the larger the number selected for cost modeling, the higher will be the calculated cost of universal service.

BCPM3: Uses total housing units. Argues that because LECs must stand ready to serve any customer requesting service, the cost of serving every potential customer location must be included in the modeling of universal service costs. To calculate per customer costs, BCPM3 divides cost of network constructed to serve all potential housing units by only the number of telephone lines the LEC actually sells.

HM 5.0: Uses households with telephones. States that universal service support should be based on the customers actually served. Per customer costs calculated by dividing cost of network constructed to serve households demanding phones by the number of telephone lines demanded by these same customers.

Discussion: According to the BCPM3 Sponsors, the 1995 national Census estimate of "Total Housing Units" is 107,879,506; the FCC Industry Analysis Division reports 100,000,000 total households and 94,000,000 households with telephones. Thus, a network modeled to serve total housing units is likely to exceed by about 8% the efficient cost of a network to serve current demand for telephone service.

In a presentation to the Commonwealth of Pennsylvania Public Utility Commission, James Dunbar of Sprint explained that BCPM's use of all housing units, including those not currently receiving telephone service causes significant increases in computed average costs. He stated that:

...the [BCPM] Model will build to all locations, all housing units, all business locations whether or not today they actually have service.

... [and that]... in some cases, for example, in the state of Maine, the entire area across the top of the state has in fact five households in it that are not served. The Model will build to those locations as if they were served. You've got to recognize therefore that if you were to try to bring in actual data, service data for loop lengths for Maine, you have over 100,000 additional feet of plant that has to be built to those five customers that had no service. That will certainly weight the average for a state.

LECs typically do not size their networks to serve locations where service is not demanded, either because housing units are not inhabited or because residents do not wish telephone service. In addition, many LEC tariffs do not obligate LECs to offer telephone service to remote locations, without payment of supplements by the customer.

In all events, the key issue is the consistency of numerator and denominator. It will surely inflate the universal service support fund to calculate per line entitlements based on the construction of a network that exceeds greatly current LEC networks in scope, yet to divide this cost only by their current scope of sales.

Score: While there may be no clear superior choice, it is clear that the BCPM3 methodology of selecting the highest possible cost numerator, and dividing it by the lowest possible lines denominator is inconsistent, and results in the largest calculated subsidy figure for the LECs.

2. Count of Residence Customer Locations by Census Block

Issue: Because the LECs do not report residence customer location counts by Census Block (CB), each model must estimate the number of customer locations to receive service in each CB.

BCPM3: Calculations are based on 1990 Census data on the total number of housing units in each CB. These 1990 CB-level counts are then inflated (or possibly reduced) by the ratio of 1995 Census estimates to 1990 Census counts by county.

HM 5.0: Calculations are based on 1997 Metromail data on the number of households in each CB. These counts are then aggregated to the CBG level. If these 1997 counts exceed Claritas 1996 CBG estimates, the Metromail counts are used for the CBs contained in this CBG. If the Claritas CBG count exceeds the Metromail CBG count, then the surplus of Claritas over Metromail is apportioned to the CB-level based on 1990 Census distributions and added to the Metromail counts.

Discussion: Data used in the BCPM3 are stale. To the extent that CBs across a county have experienced uneven growth (e.g., new housing developments established since 1990), BCPM3 will fail to reflect these counts. Because of significant changes in rural development over the last seven years, BCPM3's CB counts are likely to be inaccurate. In contrast, HM 5.0 uses 1997 data specific to the CB level when available; and when not available relies on CBG (not county) data apportioned to 1990 CB distributions only for the residual.

Score: HM 5.0 is superior to BCPM3 – both in the timeliness of its data, and in its precision by its greater use of CB/CBG-level data.

3. Count of Residence Customer Lines by Census Block

Issue: Not all residence locations demand the same number of lines. Penetration of primary and secondary residence lines must be calculated by CB.

BCPM3: BCPM3 appears to account for first and second line telephone penetration by applying flat statewide first and second line penetration factors to all CBs in the state.

HM 5.0: The number of households with at least one telephone line is estimated based on 1990 Census counts of the number of households that do not have a telephone line in their home (α). The complement of this, $(1 - \alpha)$ then represents the number of households with at least one telephone line. PNR's second line penetration model is then applied. The probability of a second line (δ) is based upon a series of logistic regressions from PNR's ReQuest™ 3 residential survey. These regressions are based upon the independent variables of age, income, and household size. The regressions are estimated conditionally upon each state and local exchange company combination.

The total number of residential access lines is thus:

$$L^* = [(1 - \alpha)(1 + \delta) * HH].$$

Finally, the modeled access line counts are normalized to reflect those numbers reported by NECA. This normalization is specific to the customer location process. The final, normalized line counts are those used in the actual customer location process. For residential data, this ratio is also used to adjust the number of residential locations.

Discussion: Because penetration rates may vary by up to 45% across different CBs in a state, the extra refinement offered in the HM 5.0 data set is important.

Score: HM 5.0 is more precise.

4. Count of Business Customer Locations and Lines by Census Block

Issue: Because the LECs do not report business customer location counts by Census Block (CB), each model must estimate the number of customer locations to receive service in each CB.

BCPM3: Uses PNR Access Line Model data – but reported only to the CB level, and not normalized to study area totals.

HM 5.0: Uses same PNR data – but reported with both CB identity and point geocodes, and normalized to study area totals.

Discussion: No dispute.

Score: Even.

5. Location of Customers Within a Census Block

Issue: CBs can be quite large – especially in rural areas. Because it is much more costly to engineer distribution plant to customers dispersed across a CB than concentrated in a single cluster, it is vital that customer locations within a CB be determined precisely to ensure that plant is modeled to match the specific configuration of customer locations.

BCPM3: Assumes that all of a CB's customer locations are distributed uniformly along the non-limited access, non-entrance/exit ramp, non-4WD, non-alley, etc. roads in the CB. "Since household and business line data are assigned at the CB level, this process requires apportioning CB line data to the corresponding microgrids. For CBs whose area is less than $\frac{1}{4}$ of a square mile, (2,640 feet by 2,640 feet), encompassing approximately three to four microgrids, household and business line data is apportioned based on the land area of the microgrid used relative to the CB's total area. For CB's with an area greater than $\frac{1}{4}$ of a square mile, household and business line data is apportioned based on relative road lengths using actual road data obtained from TIGER/Line files." (BCPM, Release 3.0 Model Methodology, December 11, 1997, Section 5.3.2, p. 26).

HM 5.0: Geocoding is used to accurately assign known customer addresses to their actual, physical locations. It involves the assignment of latitude and longitude coordinates to actual street addresses. Geocoding also allows customer location points to be associated with their CBs.

For purposes of customer location in the HM 5.0, only actual point geocodes that are returned to the 6th decimal point place of accuracy (with an accompanying CB designation) are used. The customer locations within a CB that the geocode process does not locate to this degree of accuracy

are assumed to be located uniformly along the periphery of the CB, and the geocodes implied by these placements are added to the customer location file. This addition of “surrogate” geocode points for these unlocated customers assures that the customer location file contains geocodes for 100% of the target number of customer locations by CB.

Discussion: The BCPM3 does not locate customers within a CB. Instead the BCPM3 starts with customer counts within the CB and apportions customers based on one of two methodologies: relative surface area or relative linear road distance. The first methodology, which only occurs when a CB is smaller than $\frac{1}{4}$ of a square mile, may be adequate for small CBs that have a substantial density. However, if the CB is small due to geographical or political considerations, this methodology is inadequate.

The second methodology, allocating customers based on the road lengths for CBs that are greater than $\frac{1}{4}$ of a square mile, is flawed in both concept and application. First, the BCPM3 preprocessing does not account for the differences in population clustering that often arise along different roads in its geographic areas, particularly in rural locations. Some roads will serve industrial zones, other residential areas, still others primarily retail- or service-oriented activities, and others will have a mix of one or more types – or no telephone customers at all.

The most dispositive data on whether the BCPM3’s assumption that customers are uniformly distributed along its selection of a CB’s roads is provided by geocode data of actual customer locations. Figures 1-7 show geocode maps for a selection of wire centers. Over 80% of all customer locations in these wire centers are successfully geocoded to the 6-decimal place level. The road network selected by the BCPM3 is overlaid over these geocodes. As quick inspection of the figures confirms:

- Some of BCPM3’s roads contain no geocoded customer locations;
- Some geocoded customer locations are not located on any of BCPM3’s roads;
- On the roads where customers are located, degree of customer dispersion varies widely, even within compact areas.

There are logical explanations for the failure of the BCPM3 road distribution assumptions to match reality. First, the TIGER files that BCPM3 uses for its road database are known to be incomplete. Thus, roads exist that are not in BCPM3’s TIGER. Second, many roads are connectors, or traverse unpopulated areas such as parks, swamps, mountainsides, etc. that contain no telephone customers. And finally, customer density can vary greatly within small geographies – as when a

business district abuts a residential district, or a multifamily residential development abuts a single family neighborhood.

Score: HM 5.0 customer location is clearly superior to the BCPM3 customer location for the over 70% of all customer locations are geocoded to point locations. While less desirable than point geocoding, the accuracy of the “surrogate” location method used by the HM 5.0 for the 30% or less of customer locations that are not precisely geocoded, appears to be at least as reasonable as the “road” allocation methodology used by the BCPM3 to assign 100% of its customer locations.

1. Assignment of Locations to Wire Centers

Issue: To ensure that customers are served by their present wire center, it is important that customer locations be associated with the wire center providing service to that area.

BCPM3: The BCPM, “relies on wire center data obtained from BLR ... and uses the CB level of data that falls within the corresponding wire center boundary.” (BCPM, Release 3.0 Model Methodology, December 11, 1997, Section 5.3.1, p. 24). In addition, the BCPM3 appears to use some undocumented data for wire centers for which BLR provides incomplete information, (e.g., Puerto Rico).

HM 5.0: Based on BLR data, supplemented when such data are incomplete. For example, BLR data are not available for Puerto Rico. Thus, for Puerto Rico, PNR supplements the BLR data by mapping wire centers to the municipios they serve.

Discussion: Little dispute – subject to verification that BCPM3 has performed reasonable mappings where BLR data are incomplete.

Score: Reasonably even.

2. Reasonableness and Stability of Cluster Dimensions

Issue: Clusters should conform to reasonable telephone plant engineering limitations, and these limitations should be implemented in a stable consistent fashion across the U.S. Among the important engineering restrictions are: maximum distance an analog copper loop can carry a quality signal and maximum number of lines supportable by a single DLC installation.

BCPM3: The BCPM3's "maximum ultimate grid size is typically constrained to 1/25th of a degree latitude and longitude (approximately 12,000 feet by 14,000 feet)" to "include an average maximum loop length for each CSA that is less than 12,000 feet." (BCPM, Release 3.0 Model Methodology, December 11, 1997, Section 6.2, p. 34). For practical purposes, maximum size grids that are all that BCPM3 models to in rural areas. The BCPM3 limitation of grids to 1/25th of a degree is set in its data preprocessing, and cannot be altered through user adjustment.

In the BCPM3, the preprocessing algorithm that determines serving areas, generally limits the maximum number of residence plus business lines for a single DLC to 1,000. The rationale proffered for this limitation is that "a typical DLC remote cabinet size for a large DLC, such as the "Lightspan-2000", can serve only up to 1,344 lines." (BCPM, Release 3.0 Model Methodology, December 11, 1997, Section 6.5, p. 40)

HM 5.0: The maximum permitted analog copper loop length in HM 5.0 is 18 kft – although this can be adjusted within the model by the user. This is separately ensured by both the HM 5.0 clustering process and loop engineering process. First, PNR performs a check that the distance from the centroid of a cluster to the center of the farthest location/raster within the cluster is within 18 kft, measured on a right angle basis. The threshold distance checked is actually 150 feet less than 18 kft -- to account for the maximum distance that could exist from the center of a raster to a customer located at the far vertex of the raster.

The PNR process also, generally, imposes a line limit for clusters of 1800 lines (based on an assumed maximum DLC capacity of 2,016 lines at a 90% fill level). Under certain urban conditions, it is possible for the PNR cluster process to allow clusters to form which are initially greater than 1800 lines. These oversized clusters are then subdivided by splitting evenly the cluster across its major (longer) axis. Line counts in each piece are checked against the 1800 line limit, and the splitting process continues until all resultant clusters are either under 1800 lines, or a single raster is reached. Single rasters are not divided. Within the PNR clustering algorithm, the maximum cluster size may be set by the system operator. Within the HM 5.0, maximum capacity of DLC remote terminals may be adjusted by the user.

Subsequent to the PNR clustering process, the HM 5.0 Distribution Module examines each cluster to ensure that, when modeled via a rectangular array of branch and backbone cable, maximum copper loop distances are kept to within 18,000 feet (or any other value specified by the user).

Discussion: The BCPM3 documentation states that “the design of the ultimate grids ensures that the maximum copper loop length from the DLC site to the customer for any individual customer should not exceed 18,000 feet.” (BCPM, Release 3.0 Model Methodology, December 11, 1997, Section 6.2, p. 34) Furthermore, the BCPM3’s goal appears to be to ensure that as many loops as possible not exceed a “maximum average loop length [of] 12,000 feet.” Because copper loops are capable of providing service that meets appropriate universal service criteria out to 18,000 feet, this artificial restriction ensures that the maximum BCPM3 serving area sizes will be, roughly, half of the efficient maximum serving area size. This serves to overstate costs. See Figure 8.

Perhaps more surprising is that the undefined term used by BCPM3, “average maximum loop length” appears to admit the possibility that the BCPM3 will engineer copper loops that actually exceed the 18,000 foot limit. A situation where this could occur is illustrated in Figure 9. In this figure, the greatest portion of the macrogrid’s roads are located in the lower right microgrid. Thus, even though there is a road (and thus a customer according to BCPM3’s customer location assumptions) in the upper left microgrid, the “road centroid” of the macrogrid (which the BCPM3 documentation states is the location of the DLC) is located close to its lower right corner. In the example shown, this generates a copper loop length of 20,000 feet to reach the customer in the upper left corner.

In addition to the BCPM3 maximum serving area size being unreasonably small, and possibly allowing for copper loops in excess of 18,000 feet, the size of these serving areas vary systematically across the country. Due to the curvature of the earth, the distance represented by $1/25^{\text{th}}$ of degree of latitude varies from about 1.3 miles in the northern United States (Anchorage, AK) to 2.6 miles in the southern United States (Puerto Rico) – roughly 100%. Even within the continental United States the range will be from 1.85 miles in North Dakota to 2.44 miles in Texas – about 32%. Because the BCPM3 defines grids in terms of a fixed number of degrees of latitude and longitude, the maximum serving area size in one state could be twice as large as in another state. This would result in DLC investment varying across states in similar proportion – even if customers were located in the exact same spatial configuration. Figure 10 illustrates this phenomenon in more detail. In contrast, the PNR clustering algorithm used in HM 5.0 computes distances that are invariant with respect to the location’s latitude.

The BCPM3 developers’ rationale for limiting the maximum households plus business lines for a given serving area to 1,000 is that the maximum DLC cabinet can only serve 1,344 lines. However, modern DLCs such as the DSC Litespan 2000 are designed to serve up to 2,016 lines (see,

<http://www.dsccc.com/lsp2000.htm>), and cabinets are currently available on the market that can accommodate these sizes. To restrict DLCs from being engineered to serve more than 1,000 lines limits artificially serving area size, and consequently inflates DLC costs.

Score: The BCPM3 appears not to be in compliance with the Commission's standards for universal service cost models. In particular, the BCPM3:

- Attempts to size its service areas to inefficiently small sizes;
- Fails to ensure appropriate service quality standards by permitting copper loops in excess of 18,000 feet; and
- Doesn't provide for a geographically neutral calculation of costs.

In contrast, the cluster dimensions established in the HM 5.0 assure appropriate levels of service quality, are sized to permit normal loop engineering efficiencies to be exploited, and are dimensionally stable across the entire United States.

1. Clustering of Located Customers

Issue: To economically engineer telephone distribution plant, it is vital that customers that *can* be served efficiently from a common installation (such as a DLC RT) *are* served by that installation. By locating and serving together customers that are spaced close enough to adhere to standard loop engineering specifications, these efficiencies may be ensured.

BCPM3: The BCPM does not cluster located customers, but rather uses "an unguided cookie cutter" approach of striking arbitrary maximum "cluster" boundaries at even 1/25th of a degree latitude and longitude meridian intervals. If a group of closely located customers all happen to lie within a region bounded by one of these meridian-bounded areas, they will be served together, otherwise they will be served from separate DLC sites

HM 5.0: The HM 5.0 uses a sophisticated dynamic spatial clustering algorithm to identify all customer locations that are within specified maximum distances from each other, and do not exceed specified maximum sizes.

Discussion: The HM 5.0's clustering algorithm ensures that when customers may be efficiently served together, they are. In contrast, because the BCPM3 does not base its serving area calculations on the actual location of customers, but instead determines serving areas based on

road locations and arbitrary latitude and longitude meridian, it does not do a reasonable job of clustering.

The BCPM3's "unguided cookie cutter" approach will not ensure efficient serving areas. Figure 11 provides an example of how four closely spaced customers may be served from one, two or four different DLC sites under the BCPM3 methodology – depending on where even 1/25th of a degree meridians lie.

The above concern is more than theoretical, as Figures 12-15 describing the Waterford, PA wire center demonstrate, by following the BCPM3's methodology of constraining "clusters" to lie within meridian boundaries forces an excessive number of DLC remote terminals to be engineered – even assuming BCPM3's 12,000 foot distribution length standard.

Score: Because BCPM3 doesn't really cluster customers, it cannot be considered as a reasonable model of the optimizing distribution area design decisions made by an outside plant engineer. In contrast, the HM 5.0 clustering algorithm follows faithfully normal engineering practices.

2. Calculation of Served Area Within Cluster

Issue: The amount of distribution area over which cable is laid is a key driver of OSP costs.

BCPM3: The BCPM3 calculates the area subtended by a 1000 foot swath along roads, until the total area of the grid is exhausted. This area is then assumed to exist as a square overlaying the grid's road centroid.

"For modeling purposes, the [BCPM distribution area] is a square centered about the road centroid of the distribution quadrant whose area is equal to the area encompassed by a 500 foot buffer along each side of the roads within the distribution quadrant," (BCPM, Release 3.0 Model Methodology, December 11, 1997, Section 6.7, p. 41).

HM 5.0: Computes distribution area as the area subtended by the convex hull of cluster, converted to a rectangle with equivalent aspect ratio and overlaying the cluster.

Discussion: The HM50 areas are driven directly and conservatively from the geocoded data customer location data.

The BCPM distribution plant may be substantially over- or understated, due to its “road reduction” methodology. By assuming that customers are located uniformly along roads with implied 500 foot lot depths, it is likely that computed distribution areas will be understated in rural areas where lot depths exceed 500 feet, and overstated in suburban and urban areas where lot depths fall well short of 500 feet.

In addition, the BCPM3’s calculation of distribution areas appears to be inconsistent with its driving assumption that customers are spread uniformly along roads. If this is indeed the case, why are not distribution plant cable lengths calculated simply as the total amount of road length in a grid? Even more confusing is the conflict that exists between BCPM3’s assumption here that lots have a uniform 500 foot depth, but varying frontage along roads based on density, but when calculating drop lengths, the BCPM3 assumes that lots are square.

Score: The BCPM3 suffers from an inaccurate and inconsistent methodology for computing distribution area sizes. Moreover, its methodology seems to display a rural/urban bias in implied costs. The HM 5.0 does not suffer from these deficiencies.

3. Location of Served Area Within Cluster

Issue: Determines location of SAIs and potential distribution cable lengths.

BCPM3: The BCPM3 methodology apportions the customer locations in a CB to its microgrids based on either evenly distributed across the land area or evenly distributed along the road length in each microgrid. The BCPM3 then determines an ultimate grid or serving area that may be comprised of 64 microgrids containing 16 microgrids in each quadrant. The BCPM3 then calculates the total area for a quadrant (which may contain up to 16 microgrids), based on an arbitrary 500 foot buffer on either side of the road with customers evenly distributed along them. The BCPM3 then moves the customers that were originally placed along roads in a microgrid and places them in the middle of the quadrant evenly distributed over the road buffered areas (which has now been squared).

HM 5.0: Places the distribution area over the determined location of customers. DLC sites are located within this distribution area.

Discussion: The BCPM3 moves customers from the microgrid where their “road” was located, to a distribution area that may be remote from that microgrid – and may contain no customers or roads. DLCs are placed at the “road centroid,” a location where there may also be no customers or roads. It is

difficult to understand why this methodology should yield an accurate assessment of distribution costs.

Score: Because the BCPM3 models distribution plant and DLCs at locations where these items were not “located,” it is almost certainly inferior to the HM 5.0 methodology that places these items where the customers were actually located.